arXiv:1411.1041v2 [math.NT] 25 Nov 2014

CANONICAL HEIGHTS FOR CORRESPONDENCES

PATRICK INGRAM

Abstract. The canonical height associated to a polarized endomporhism of a projective variety, constructed by Call and Silverman and generalizing the Néron-Tate height on a polarized Abelian variety, plays an important role in the arithmetic theory of dynamical systems. We generalize this construction to polarized correspondences, prove various fundamental properties, and show how the global canonical height decomposes as an integral of a local height over the space of absolute values on the algebraic closure of the field of definition.

1. Introduction

1.1. **Dynamics of correspondences.** Let X be an algebraic variety defined over some field K, and let C be a *correspondence* on X, by which we mean a subvariety $C \subseteq X \times X$ such that the coordinate projections are both finite and surjective. We may think of the points $X(\overline{K})$ as a set of vertices, with the points $C(\overline{K})$ defining directed edges between those vertices, and inquire about the structure of the resulting graph. For instance, if $f : X \to X$ is a finite morphism, then the graph of f defines a correspondence on X, and the paths in the aforementioned directed graph correspond to orbits in the sense of dynamics. Thus the study of paths in the directed graph associated to a correspondence is a natural generalization of the study of dynamical systems.

The purpose of this note is to associate a canonical height to certain correspondences over number fields in a way which generalizes the Call-Silverman canonical height associated to a polarized algebraic dynamical system [8]. We note that Autissier [1, 2] has previously considered canonical heights associated to correspondences, but our construction differs significantly from his; the relation between the two is discussed in Remark 1.2 below.

In the study of the iteration of a morphism $f : X \to X$, orbits are defined entirely in terms of their initial points (we shall henceforth refer to this as the *deterministic* case), and so many statements about orbits end up being quantified over the set of points $X(K)$ or $X(\overline{K})$, rather than over the set of orbits *per se*. In the study of correspondences this becomes inappropriate, and so we define below a *path space* $\pi : \mathscr{P} \to X$ for a given correspondence, parametrizing paths through the aforementioned directed graph, with initial vertex indicated by π. We will also construct a *shift map* $\sigma : \mathscr{P} \to \mathscr{P}$ which corresponds to stepping forward once along the path (or, equivalently, forgetting the first vertex of a path). These objects can be constructed in the category of K-schemes, but generally not in the category of varieties over K, a distinction which becomes important in our discussion of local and global heights. We define $\mathscr{P}$ and σ explicitly in Section 2, but for now

Date: November 26, 2014.

the reader in encouraged to simply think of $\mathscr{P}(K)$ as being the set of paths whose points all sit in $X(K)$.

1.2. **The global canonical height.** When K is a number field, one of the central tools in the study of polarized algebraic dynamical systems over K is the canonical height $\hat{h}_{X,f,L} : X(\overline{K}) \to \mathbb{R}$ of Call and Silverman [8], a construction extending classical work of Néron and Tate in the context of abelian varieties. Here X is a projective variety, $f : X \to X$ is an endomorphism, and L is an ample divisor on X satisfying $f^*L \sim \alpha L$ in $\mathrm{Div}(X) \otimes \mathbb{R}$, for some real $\alpha > 1$. Our first result is the existence of a generalization of the Call-Silverman canonical height in the context of dynamical correspondences.

Let C be a correspondence on X as above, and let $x, y : C \to X$ be the coordinate projections. We will say that this correspondence is *polarized* if and only if there exists an ample Cartier divisor $L \in \mathrm{Div}(X) \otimes \mathbb{R}$ and a real number $\alpha > 1$ such that y^*L is linearly equivalent to αx^*L. It is easy to check that in the case where C is the graph of a morphism $f : X \to X$, the correspondence is polarized just in case the algebraic dynamical system is polarized in the usual sense.

Theorem 1.1. *Given a correspondence C on X, defined over a number field K and polarized by $L \in \mathrm{Div}(X) \otimes \mathbb{R}$ and $\alpha > 1$, there exists a function $\hat{h}_{X,C,L} : \mathscr{P}(\overline{K}) \to \mathbb{R}$ such that for all $P \in \mathscr{P}(\overline{K})$*

$$\hat{h}_{X,C,L}(\sigma(P)) = \alpha \hat{h}_{X,C,L}(P)$$

and

$$\hat{h}_{X,C,L}(P) = h_{X,L}(\pi(P)) + O(1).$$

As one might suspect, in the case where C is the graph of a morphism $f : X \to X$, we simply recover the usual canonical height (and in this case $\pi : \mathscr{P} \to X$ turns out to be an isomorphism).

Remark 1.2. Autissier [2] constructs a canonical height associated to a correspondence under certain conditions, and we reflect here on the relation between the constructions. Specifically, suppose that we are in the context of Theorem 1.1 with the additional hypothesis that $\alpha \in \mathbb{Q}$. Then Autissier constructs a real-valued function $h'_{X,C,L}$ on 1-cycles of X with the property that

$$h'_{X,C,L}(C_*D) = \deg(x) h'_{X,C,L}(D) + \deg(x) \deg(D_K) h_0,$$

where $C_* = y_* x^*$ and h_0 is the *charge* of the correspondence, defined in [2].

Autissier's heigh is defined on (1-cycles on) $X(\overline{K})$, while the height constructed in Theorem 1.1 is a function on $\mathscr{P}(\overline{K})$. Note, however, that for any $a \in X(\overline{K})$, the fibre $\pi^{-1}(a) \subseteq \mathscr{P}(\overline{K})$ naturally admits the structure of a probability space. One may construct this measure as the inverse limit of the uniform distribution on paths of length n (weighting the paths appropriately when there is ramification), and it is easy to show that $\hat{h}_{X,C,L}$ becomes a measurable function on this space, that is, a random variable. Let $\mathbb{E}\hat{h}_{X,C,L}(a)$ denote the expected value of $\hat{h}_{X,C,L}(P)$ for $\pi(P) = a$, in other words

$$\mathbb{E}\hat{h}_{X,C,L}(a) = \int_{\pi(P)=a} \hat{h}_{X,C,L}(P) d\mu(P),$$

where μ is the distribution defined above. It is easy to confirm from Theorem 1.1 that (extending this function linearly to 1-cycles)

$$\mathbb{E}\hat{h}_{X,C,L}(C_*a) = \deg(x)\mathbb{E}\hat{h}_{X,C,L}(a).$$

Since

$$\mathbb{E}\hat{h}_{X,C,L}(a) = h_{X,L}(a) + O(1) = h'_{X,C,L}(a) + O(1),$$

we see that in the special case $h_0 = 0$ and $\deg(x) > 1$, we must actually have $h'_{X,C,L} = \mathbb{E}\hat{h}_{X,C,L}$. In other words, at least in some cases, Autissier's canonical height is the average value of the canonical height constructed in this note.

One of the useful features of the canonical height in deterministic arithmetic dynamics is that its vanishing identifies the preperiodic orbits. A path $P \in \mathscr{P}(\overline{K})$ will be called *repetitive* if and only if $\pi \circ \sigma^n(P) = \pi \circ \sigma^m(P)$ for some $n \neq m$ (in other words, if and only if the path passes through some vertex twice). Note that, unlike in the case of deterministic dynamics, repetitive paths need not repeat more than once.

Corollary 1.3. *The set of points $a \in X(\overline{K})$ such that $\pi^{-1}(a) \subseteq \mathscr{P}(\overline{K})$ contains a repetitive path is a set of bounded height. In particular, there are only finitely many $a \in X(K)$ which are the initial vertex of some repetitive path over $\overline{K}$.*

Note that it is quite important, in the statement of the corollary, that we consider paths defined over the algebraic closure. Unlike in the deterministic case, paths are not generally defined over the field of definition of the initial point, and indeed are typically not even defined over a finite extension of this field. Furthermore, the finiteness in the above statement must apply to initial points, rather than paths, since it is easy to construct examples wherein $\pi^{-1}(a)$ contains uncountably many repetitive paths for a single $a \in X(K)$. In order to elucidate the situation, we present an example below. A path $P \in \mathscr{P}(\overline{K})$ is said to be *finitely supported* if and only if the set of points of the form $\pi \circ \sigma^n(P) \in X(\overline{K})$ is finite as n varies, in other words, if and only if there are only finitely many distinct points of X occurring as vertices in the path. Note that this is a stronger condition than being repetitive, and is strong enough to immediately imply $\hat{h}_{X,C,L}(P) = 0$.

Example 1.4. Consider the correspondence defined on $\mathbb{P}^1$ by the closure in $\mathbb{P}^1 \times \mathbb{P}^1$ of the curve

$$C : y^2 = x^3 - x + 1,$$

with x and y corresponding to the coordinate projections. Here $\mathscr{P}(\mathbb{Q})$ turns out to be uncountable, but it is also true that every path in $\mathscr{P}(\mathbb{Q})$ is finitely supported. There are two ways to see this latter claim.

First of all, the curve

$$\left\{ \begin{array}{c} y^2 = x^3 - x + 1 \\ z^2 = y^3 - y + 1 \end{array} \right\}$$

parametrizing paths of length 3 under this correspondence is a curve of genus 6. Hence, by Faltings' Theorem, there are only finitely many points $a \in \mathbb{P}^1(\mathbb{Q})$ from which we may build a path of length 3, let alone an infinite path.

From another view, suppose that $K/\mathbb{Q}$ is the field of fractions of a DVR with valuation v, and that $(a, b) \in C(K)$, with $v(a) < 0$. Then we have $v(b) = \frac{3}{2}v(a)$ and so, in particular, there can be no path $P \in \mathscr{P}(K)$ with initial vertex $\pi(P) = a$ (such a path would contradict v being discrete). Turning this around, if $P \in \mathscr{P}(\mathbb{Q})$,

then each edge $a \to b$ of P must correspond to an integral point $(a, b) \in C(\mathbb{Q})$. It is an easy exercise to show that the integral points on $C \subseteq \mathbb{P}^1 \times \mathbb{P}^1$ are precisely

$$\{(-1, \pm 1), (0, \pm 1), (1, \pm 1), (3, \pm 5), (5, \pm 11), (56, \pm 419), (\infty, \infty)\}.$$

In particular, the paths defined over $\mathbb{Q}$ are exactly those sequences with initial vertex 0, 1, or -1, and subsequent vertices ± 1. This shows that $\mathscr{P}(\mathbb{Q})$ is uncountable, but is contained in only three fibres of $\pi : \mathscr{P} \to \mathbb{P}^1$, and consists entirely of finitely supported paths.

It is not difficult to extend this argument from integrality to show that, for many interesting correspondences over a number field K, the collection $\mathscr{P}(K)$ of K-rational paths is rather small, or at least rather pathological. Indeed, Example 1.4 can easily be generalized to a large class correpondences of the form $g(y) = f(x)$, with g and f polynomials, via results of Bilu and Tichy [7] and Avanzi and Zannier [3]. The aim of this remark is simply to convince the reader that one should really consider paths defined over $\overline{K}$, and that these are not *a priori* defined over any finite extension of the base (see below).

1.3. **Specialization.** Our next result is a specialization theorem along the lines of Theorem 4.1 of [8], but again the setting requires us to introduce some additional technical details. If K is a number field and U/K is a curve, then we may consider a family of correspondences over U, that is a family $X \to U$ and $C \subseteq X \times_U X$ where, as usual, the projection maps are finite and surjective. Such an object is said to be *polarized* if and only if there is an ample $L \in \mathrm{Div}(X) \otimes \mathbb{R}$ and a real number $\alpha > 1$ such that $y^*L - \alpha x^*L$ is linearly equivalent to a divisor fibral to the projection $C \to U$. One checks, then, that for all but finitely many $t \in U(\overline{K})$, the fibres X_t, C_t, L_t give a polarized correspondence over $K(t)$.

Writing $t = \eta$ for the generic point of U, Theorem 4.1 of [8] gives for deterministic arithmetic dynamical systems the estimate

$$\hat{h}_{X,f,L}(P_t) = \left(\hat{h}_{X_\eta, f_\eta, L_\eta}(P_\eta) + o(1)\right) h_U(t)$$

where h_U is any height on U with respect to a divisor class of degree 1, and $o(1) \to 0$ as $h_U(t) \to \infty$ (see [11, 14, 15, 16] for improvements and variations in some special cases).

Just as one would not like to consider only paths defined over a number field, however, it turns out that considering only paths defined over U, or finite covers thereof, is overly restrictive; it is easy to construct schemes $\mathscr{U} \to U$ such that $\mathscr{P}(\mathscr{U})$ contains many interesting points which do not factor through finite $V \to U$. In simpler terms, elements of $\mathscr{P}(U)$ correspond to paths whose nodes are all $k(U)$-rational and, since function fields of varieties are generally not algebraically closed, this restriction rules out many interesting paths.

In order to give a wider class of paths to which our theorem will apply, we allow ourselves to consider $\mathscr{U}$-valued paths, for arbitrary $\mathscr{U} \to U$, as long as the initial vertex is a U-point.

Theorem 1.5. *Fix a family of polarized correspondences over U, as above, a dominant morphism of schemes $\varphi : \mathscr{U} \to U$, and a path $P \in \mathscr{P}(\mathscr{U})$ with $\pi(P) \in X(U) \subseteq X(\mathscr{U})$. Then for any height h_U on U relative to a divisor class of degree one,*

$$\hat{h}_{X_{\varphi(t)}, C_{\varphi(t)}, L_{\varphi(t)}}(P_t) = \left(\hat{h}_{X_{\varphi(\eta)}, C_{\varphi(\eta)}, L_{\varphi(\eta)}}(P_\eta) + o(1)\right) h_U(\varphi(t)),$$

where $o(1) \to 0$ *as* $h_U(\varphi(t)) \to \infty$. *In particular, the height on the generic fibre may be computed as*

$$\hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta) = \lim_{h_U(\varphi(t))\to\infty} \frac{\hat{h}_{X_{\varphi(t)},C_{\varphi(t)},L_{\varphi(t)}}(P_t)}{h_U(\varphi(t))}.$$

Corollary 1.6. *With the assumptions of Theorem 4.1, let*

$$Z = \{t \in \mathscr{U}(\overline{K}) : P_t \text{ is finitely supported}\}.$$

Then either $\varphi(Z) \cap U(K)$ *is finite, or else* $\hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta) = 0$.

1.4. **Local canonical heights.** The construction of the canonical height above parallels that of Call and Silverman very closely, but in the local theory we see a more significant divergence. One of the fundamental facts about varieties X over fields K is that

$$X(\overline{K}) = \bigcup_{[L:K]<\infty} X(L). \tag{1}$$

Indeed, this is true whenever X is a K-scheme of finite type. This means that in the development of local heights for varieties, it is always possible to work over number fields; the theory over the algebraic closure follows once one shows that the values of the functions are independent of the particular extension.

Since $\mathscr{P}$ is not of finite type, (1) does not hold. Indeed, Example 1.4 and its generalizations proffer a bounty of instances in which, for every finite extension L/K, the paths in $\mathscr{P}(L)$ consist of fairly pathological examples, far from the generic behaviour in $\mathscr{P}(\overline{K})$. To remedy this, we employ Gubler's theory of M-fields. The following is essentially present in [12].

Theorem 1.7. *Let* k *be the algebraic closure of a number field, and let* M_k *denote the usual set of places of* k. *Then there exists a* σ*-algebra* $\mathscr{F}$ *of subsets of* M_k *and a* σ*-finite measure* $\mu : \mathscr{F} \to \mathbb{R}$ *such that the following holds: For any finite* L/K *and any place* $v \in M_L$, *the set*

$$A_{v,L} = \{w \in M_k : v \text{ is the restriction to } L \text{ of } w\}$$

is measurable, and

$$\mu(A_{v,L}) = \frac{[L_v : K_v]}{[L : K]}. \tag{2}$$

The ratio in (2) appears frequently in weighted sums in the theory of heights, and one should think of integrating with respect to this measure as a version of these weighted sums. In particular, for any non-zero $\alpha \in k$, the functions $v \mapsto \log|\alpha|_v$ and $v \mapsto \log^+|\alpha|_v$ are both measurable, and satisfy

$$\int_{M_k} \log|\alpha|_v d\mu(v) = 0$$

(essentially the product formula) and

$$\int_{M_k} \log^+|\alpha|_v d\mu(v) = h(\alpha).$$

This construction allows us to present a theory of local heights for correspondences.

Theorem 1.8. *Let C be a correspondence on X, polarized by an ample class L and a real $\alpha > 1$, let $\mathscr{W} = \mathscr{P} \setminus \pi^*(L)$, and fix a local height $\lambda_{X,L}$ on X relative to L. Then there exists a function*

$$\hat{\lambda}_{X,C,L} : \mathscr{W} \times M_{\overline{K}} \to \mathbb{R}$$

such that the following hold:

(1) $\hat{\lambda}_{X,C,L}(\cdot, v) : \mathscr{W}(\overline{K}) \to \mathbb{R}$ *is v-adically continuous.*
(2)
$$\hat{\lambda}_{X,C,L}(P, v) = \lambda_{X,L}\left(\pi(P), v\right) + O(1),$$
where the implied constant depends on the place.
(3) *For some $f \in k(C) \otimes \mathbb{R}$ satisfying $y^*L - \alpha x^*L = \mathrm{div}(f)$, we have*
$$\hat{\lambda}_{X,C,L}(\sigma(P), v) = \alpha \hat{\lambda}_{X,C,L}(P, v) + \log |f \circ \varepsilon(P)|_v,$$
where $\varepsilon : \mathscr{P} \to C$ takes a path to the point on C which gives the first edge.
(4) *For $P \in \mathscr{P}(\overline{K})$ function $\hat{\lambda}_{X,C,L}(P, \cdot)$ is measurable, and*
$$\hat{h}_{X,C,L}(P) = \int_{M_{\overline{K}}} \hat{\lambda}_{X,C,L}(P, v) d\mu(v).$$

1.5. **Questions.** Before proceeding with a more formal development, we pose a few natural questions arising from the results in this paper. For this discussion, we fix a number field K, a projective variety X (which, for the sake of simplicity, we will assume to be $\mathbb{P}^N$), and a correspondence C on X.

Question 1.9. If the correspondence C is polarized by L, for which $P \in \mathscr{P}(\overline{K})$ do we have $\hat{h}_{X,C,L}(P) = 0$?

In the case of deterministic arithmetic dynamics, this is fairly well understood, and even in the present context there are some obvious sufficient conditions. But we also produce a class of examples in Section 3 which contradict what might be an initial naive answer to Question 1.9.

In order to pose the next question, we note that it follows from Corollary 1.3 that there exists an B such that, if $P \in \mathscr{P}(\overline{K})$ is preperiodic, and $\pi(P) \in X(K)$, then there is a path $Q \in \mathscr{P}(\overline{K})$ with $\pi(P) = \pi(Q)$, and $\sigma^n(Q) = \sigma^m(Q)$ for some $0 \leq m < n \leq B$. In other words, a preperiodic path whose initial vertex is K-rational must share this initial vertex with a preperiodic path of bounded length.

Question 1.10. How uniform can we make the value B in the above remark? Is there a value which depends only on the degrees of the coordinate projections, and the degree of K (in the case $X = \mathbb{P}^N$)? Note that this is a generalization of the uniform boundedness conjecture of Morton and Silverman [20].

2. Geometry of correspondences

In this section we define some notions needed in the geometry of a correspondence. We may proceed with the following constructions in any category admitting fibre products and inverse limits, but in order to keep things tangible we work in the category of S-schemes, for some fixed scheme S. It is tempting to work in the category of varieties over a field, but this too restrictive to allow a satisfactory construction of the path space. The reader who is not interested in the content of Section 4 could take S to be the spectrum of a field.

Let X be a separated, integral S-scheme of finite type (i.e., a variety if S is the spectrum of a field), and let $C \subseteq X \times_S X$ be a subscheme such that the two projection maps are finite and surjective. Our first goal is to construct the path space $\mathscr{P}_C$ (we omit the subscript when the context makes it obvious) of the correspondence on X defined by these data. To do this, we construct a sequence of S-schemes parametrizing paths of finite length, and take $\mathscr{P}$ to be their inverse limit. There is then a natural interpretation of the correspondence as a dynamical system on $\mathscr{P}$. In particular, there is a *shift map* $\sigma : \mathscr{P} \to \mathscr{P}$ which corresponds to deleting the first vertex in a path.

Theorem 2.1. *Let C be a correspondence on the S-scheme X. Then there exists a separated, integral S-scheme $\mathscr{P}$, surjective morphisms $\pi : \mathscr{P} \to X$, $\varepsilon : \mathscr{P} \to C$, and a finite surjective morphism $\sigma : \mathscr{P} \to \mathscr{P}$ making the following diagram commute.*

$$\begin{array}{ccccc} \mathscr{P} & & \xrightarrow{\sigma} & & \mathscr{P} \\ {\scriptstyle\pi}\downarrow & \searrow^{\varepsilon} & & & \downarrow{\scriptstyle\pi} \\ X & \xleftarrow[x]{} & C & \xrightarrow[y]{} & X \end{array}$$

For a given path $P \in \mathscr{P}$, we should think of $\pi(P) \in X$ as the initial vertex of P, $\varepsilon(P) \in C$ as the initial edge, and $\sigma(P) \in \mathscr{P}$ as the path obtained by forgetting both of these (or shifting along the path). Thus, the commutative diagram in the statement of the theorem says roughly that the terminal vertex of the initial edge of a path is the initial vertex of the path shifted once (and that the initial vertex of the initial edge is also the initial vertex of the entire path).

Proof of Theorem 2.1. We define a sequence of S-schemes $\mathscr{P}_n$, along with morphisms $\pi_n, \tau_n : \mathscr{P}_n \to X$, as follows. First, we take $\mathscr{P}_1 = X$, with $\pi_1 = \tau_1$ the identity map. Now, for each n, we define $\mathscr{P}_{n+1}$ to be the base extension

$$\begin{array}{ccc} \mathscr{P}_{n+1} & \xrightarrow{\beta_n} & C \\ {\scriptstyle\alpha_n}\downarrow & & \downarrow{\scriptstyle x} \\ \mathscr{P}_n & \xrightarrow[\tau_n]{} & X \end{array}$$

The maps π_{n+1}, τ_{n+1} are defined by $\tau_{n+1} = y \circ \beta_n$, and $\pi_{n+1} = \pi_n \circ \alpha_n$.

Intuitively, $\mathscr{P}_n$ parametrizes paths of length n, with initial and terminal vertices given by π_n and τ_n, respectively. By the fibre product construction, $\mathscr{P}_{n+1}$ parametrizes pairs consisting of a path of length n and an edge, such that the terminal vertex of the path matches the initial vertex of the edge; in other words, a path of length $n+1$. Note that, by the standard properties of fibre products, each $\mathscr{P}_n$ is a separated, integral scheme of finite type over S.

Now consider the inverse system of S-schemes

$$\cdots \xrightarrow{\alpha_{n+1}} \mathscr{P}_{n+1} \xrightarrow{\alpha_n} \cdots \xrightarrow{\alpha_2} \mathscr{P}_2 \xrightarrow{\alpha_1} \mathscr{P}_1,$$

and let $\mathscr{P} = \varprojlim \mathscr{P}_n$ (see [18, Chapter 31]). Note that $\mathscr{P}$ is not, in general, of finite type over S. That $\mathscr{P}$ is a separated, integral S-scheme follows simply from the fact that these properties are preserved under base extension and taking limits.

The inverse limit construction gives morphisms $\mathscr{P} \to \mathscr{P}_n$ for every n, which are surjective and commute with the transition maps α_n. Since $\mathscr{P}_1 \cong X$ and $\mathscr{P}_2 \cong C$,

we thus have surjective maps $\pi : \mathscr{P} \to X$ and $\varepsilon : \mathscr{P} \to C$ which satisfy $\pi = x \circ \varepsilon$ (the left-hand-side of the claimed commutative diagram).

It remains to construct $\sigma : \mathscr{P} \to \mathscr{P}$. Note that, by diagram chasing, we have $C \times_X \mathscr{P}_n \cong \mathscr{P}_{n+1}$, with respect to the maps $y : C \to X$ and $\pi_n : \mathscr{P}_n \to X$. In other words, paths of length $n+1$ might just as well be constructed from paths of length n by adding an edge at the beginning. Note that the resulting map $\mathscr{P}_{n+1} \to C$ corresponds to $\alpha_2 \circ \cdots \circ \alpha_n$ under the association $C \cong \mathscr{P}_2$. Taking limits, this gives (by [18, Chapter 39, Lemma 2.3])

$$C \times_X \mathscr{P} = C \times_X \varprojlim \mathscr{P}_n = \varprojlim (C \times_X \mathscr{P}_n) \cong \varprojlim \mathscr{P}_{n+1} = \mathscr{P},$$

where the latter is viewed as an X-scheme by $\pi : \mathscr{P} \to X$, and C is viewed as an X-scheme by $y : C \to X$. Furthermore, the implied map $C \times_X \mathscr{P} \to C$ coincides, from the construction, with $\varepsilon : \mathscr{P} \to C$. In other words, we have a base extension diagram

$$\begin{array}{ccc} \mathscr{P} & \longrightarrow & \mathscr{P} \\ {\scriptstyle\varepsilon}\downarrow & & \downarrow{\scriptstyle\pi} \\ C & \xrightarrow[y]{} & X, \end{array}$$

and we call the top morphism σ. The fact that $\sigma : \mathscr{P} \to \mathscr{P}$ is finite, separated, and surjective follows from it being a base extension of a finite, surjective morphism of varieties.

□

The proof of the finiteness of $\sigma : \mathscr{P} \to \mathscr{P}$ belies the simplicity of the motivation. Essentially, given a path Q there should exist only finitely many paths P with $\sigma(P) = Q$ simply because there are only finitely many points on X which could be the initial vertex of a path whose second vertex is $\pi(Q)$. This follows just because the two projection maps $C \to X$ are finite.

In deterministic dynamics, the notion of preperiodicity is very important. There are numerous possible generalizations of this to the dynamics of correspondences, and we name three.

Definition 2.2. Let X be a projective variety, let C be a correspondence on X, let $\mathscr{P}$ be the associated path space, and let $P \in \mathscr{P}$. We say that P is

(1) *preperiodic* if and only if there exist $n \neq m$ with $\sigma^n(P) = \sigma^m(P)$;
(2) *finitely supported* if and only if the quantity $\pi \circ \sigma^n(P)$ takes only finitely many values as n varies; and
(3) *repetitive* if and only if there exist $n \neq m$ with $\pi \circ \sigma^n(P) = \pi \circ \sigma^m(P)$.

We say that $a \in X$ is *constrained* relative to the correspondence C if and only if there exists a repetitive path $P \in \pi^{-1}(a)$.

Remark 2.3. It is easy to see that every preperiodic path is finitely supported, and every finitely supported path is repetitive, but in general the three conditions are not equivalent (as they are in the deterministic case). Example 1.4 contains paths which are finitely supported but not preperiodic, while Example 3.1 below gives a path which is repetitive but not finitely supported. Note, however, that the condition that there exists a repetitive path P with $\pi(P) = a$ *is* equivalent to the condition that there exists a preperiodic (and hence finitely supported) path P with $\pi(P) = a$.

Remark 2.4. For each $n \geq 2$, let $\mathscr{P}_n$ be the S-scheme parametrizing paths of length n as above, with initial and terminal vertices given by the maps $\pi_n, \tau_n : \mathscr{P}_n \to X$. This gives a morphism $(\pi_n, \tau_n) : \mathscr{P}_n \to X \times_S X$. The points of X which land in the image of this morphism under the diagonal embedding $X \to X \times_S X$ are precisely the initial vertices of cyclic paths of length $n-1$. Since we should expect $\mathscr{P}_n$ to have the same dimension as X, one should typically expect this intersection to be 0-dimensional, but additional hypotheses are needed to make this concrete.

3. The global canonical height

Let X be a projective variety over K, and let C be a correspondence over X. Recall that C is *polarized* if and only if there is an ample (Cartier) divisor $L \in \mathrm{Div}(X) \otimes \mathbb{R}$ and a real number $\alpha > 1$ such that $y^*L \sim \alpha x^*L$, where $x, y : C \to X$ are the coordinate projections. In this section we will prove Theorem 1.1.

Before proceeding with the proof, we note that nowhere do we use the condition that L is ample. This condition arises in Corollary 1.3, and is likely to play a role in the determination of where the canonical height vanishes.

Proof of Theorem 1.1. Let $P \in \mathscr{P}(\overline{K})$. By the standard properties of Weil heights [13, Theorem B.3.2, p. 184], and the commutative diagram in Proposition 2.1, we have

$$\begin{aligned} h_{X,L}\left(\pi \circ \sigma(P)\right) &= h_{X,L}(y \circ \varepsilon(P)) \\ &= h_{C,y^*L}(\varepsilon(P)) + O(1) \\ &= \alpha h_{C,x^*L}(\varepsilon(P)) + O(1) \\ &= \alpha h_{X,L}(x \circ \varepsilon(P)) + O(1) \\ &= \alpha h_{X,L}(\pi(P)) + O(1). \end{aligned}$$

(Note in particular, that although P is not a point on a projective variety, both $\pi(P)$ and $\varepsilon(P)$ are, and so we are not applying the height machine out of context.) Thus it follows that

$$\left|\alpha^{-1} h_{X,L}(\pi \circ \sigma(P)) - h_{X,L}(\pi(P))\right| \leq \kappa$$

for some constant $\kappa \geq 0$, for all $P \in \mathscr{P}(\overline{K})$. We now apply a standard telescoping sum argument, due to Tate (seen also in [8]). By the triangle inequality, we have

(3)
$$\left|\alpha^{-n} h_{X,L}(\pi \circ \sigma^n(P)) - h_{X,L}(\pi(P))\right| \leq \left(1 + \alpha^{-1} + \cdots + \alpha^{-n+1}\right)\kappa \leq \left(\frac{\alpha}{\alpha - 1}\right)\kappa.$$

Applying this with $\sigma^m(P)$ in place of P, we have

$$\left|\alpha^{-(m+n)} h_{X,L}(\pi \circ \sigma^{m+n}(P)) - \alpha^{-m} h_{X,L}(\pi \circ \sigma^m(P))\right| \leq \alpha^{-m}\left(\frac{\alpha}{\alpha - 1}\right)\kappa.$$

This shows that the sequence $\alpha^{-n} h_{X,L}(\pi \circ \sigma^n(P))$ of real numbers is Cauchy, whereupon the limit

$$\hat{h}_{X,C,L}(P) = \lim_{n \to \infty} \alpha^{-n} h_{X,L}(\pi \circ \sigma^n(P))$$

exists. The first property in the theorem follows immediately from the definition, while the second follows by taking the limit of (3) as $n \to \infty$. □

It is easy to see from the definition that $\hat{h}_{X,C,L}(P) = 0$ when P is finitely supported (and so when P is preperiodic), although this is not generally the case when P is merely repetitive, as the following example shows.

Example 3.1. Let $C \subseteq \mathbb{P}^1 \times \mathbb{P}^1$ be the correspondence defined by (the closure of) the curve $y^2 = x^3 + 1$. This correspondence is polarized, with L any ample divisor on $\mathbb{P}^1$ and $\alpha = 3/2$.

It is not hard to check, by making explicit the bound on the difference between the canonical height and the Weil height, that any path P which begins as

$$P : 0 \to 1 \to 2^{1/2} \to (2^{3/2} + 1)^{1/2} \to \cdots$$

satisfies $\hat{h}_{X,C,L}(P) > 0$. But now note that the repetitive path

$$Q : 0 \to -1 \to 0 \to 1 \to 2^{1/2} \to \cdots$$

satisfying $\sigma^2(Q) = P$ must have $\hat{h}_{X,C,L}(Q) = \frac{4}{9}\hat{h}_{X,C,L}(P) > 0$. Indeed, we may use this example to produce paths $Q \in \mathscr{P}(\overline{\mathbb{Q}})$ of arbitrarily small positive canonical height with $\pi(Q) \in \mathbb{P}^1(\mathbb{Q})$ (indeed, with $\pi(Q) = 0$). This is a phenomenon not present in deterministic arithmetic dynamics.

Remark 3.2. À propos of the discussion that many interesting $\overline{K}$-points of $\mathscr{P}$ are not defined over any number field, it is worth noting that a path $P \in \mathscr{P}(\overline{K})$ is finitely supported if and only if we have both $\hat{h}_{X,C,L}(P) = 0$ and $P \in \mathscr{P}(E)$ for some finite extension E/K. The question of when the vanishing of the canonical height implies finite support of a path is thus intimately linked to the question of which paths are defined over some number field. Note that if $\hat{h}_{X,C,L}(P) = 0$ but P is not E-rational for any finite E/K, then the points $\pi \circ \sigma^n(P)$ have unbounded algebraic degree, and height zero with respect to a certain Weil height on X; these are precisely the antecedent conditions of many equidistribution results.

In general, it is not clear what conditions are necessary for $\hat{h}_{X,C,L}(P) = 0$, even over number fields (as opposed to functions fields, in which context we expect the answer to be even more subtle [4, 6]). For instance, let $X = \mathbb{P}^1$ and let $C \subseteq \mathbb{P}^1 \times \mathbb{P}^1$ be the correspondence defined by the closure of $y^2 = x^3$. Then any $P \in \mathscr{P}(\overline{\mathbb{Q}})$ with $\pi(P)$ a root of unity satisfies $\hat{h}_{X,C,L}(P) = 0$, but there are certainly non-repeating paths of this form.

More generally, if $f, g : X \to X$ are two finite morphisms whose set of preperiodic points coincide, then the correspondence defined on X by $g(y) = f(x)$ admits non-repeating paths of canonical height zero. In light of characterizations such as those in [5, Theorem 1.2], it would be of interest to determine whether or not these are the only sorts of examples.

Before concluding this section, we note that there is a natural topology on $\mathscr{P}(\overline{K})$, which we will call the *tree topology*. In particular, equip $\mathscr{P}_n(\overline{K})$ with the discrete topology for all n, and define the tree topology on $\mathscr{P}_n(\overline{K})$ to be the weakest which makes each of the projection maps $\mathscr{P}(\overline{K}) \to \mathscr{P}_n(\overline{K})$ continuous. A base for this topology is simply the collection of sets

$$U_{P,n} = \left\{ Q \in \mathscr{P}(\overline{K}) : \pi \circ \sigma^m(Q) = \pi \circ \sigma^m(P) \text{ for all } 0 \leq m \leq n \right\},$$

for $n \geq 0$ and $P \in \mathscr{P}(\overline{K})$. It is natural to ask how the canonical height behaves with respect to this topology.

Proposition 3.3. *The function $\hat{h}_{X,C,L} : \mathscr{P}(\overline{K}) \to \mathbb{R}$ is continuous with respect to the tree topology, and each fibre $\mathscr{P}_a(\overline{K}) = \pi^{-1}(a) \subseteq \mathscr{P}(\overline{K})$ is compact.*

Proof. The compactness of $\mathscr{P}_a(\overline{K})$ can be obtained as a general fact about inverse limits of compact spaces, since $\mathscr{P}_a(\overline{K})$ is the inverse limit of $\mathscr{P}_{a,n}(\overline{K})$, a finite discrete space for each n.

Note that we have constructed, above, a constant $\kappa' \geq 0$ such that for any $P \in \mathscr{P}(\overline{K})$, we have

$$\left|\hat{h}_{X,C,L}(P) - h_{X,L}(\pi(P))\right| \leq \kappa'.$$

Now, supposing that $Q \in U_{P,n}$, we have

$$\begin{aligned}\left|\hat{h}_{X,C,L}(P) - \hat{h}_{X,C,L}(Q)\right| &= \alpha^{-n}\left|\hat{h}_{X,C,L}(\sigma^n(P)) - \hat{h}_{X,C,L}(\sigma^n(Q))\right| \\ &\leq \alpha^{-n}\left|h_{X,L}(\pi\circ\sigma^n(P)) - h_{X,L}(\pi\circ\sigma^n(Q))\right| \\ &\quad +\alpha^{-n}\left|\hat{h}_{X,C,L}(\sigma^n(P)) - h_{X,L}(\pi\circ\sigma^n(P))\right| \\ &\quad +\alpha^{-n}\left|\hat{h}_{X,C,L}(\sigma^n(Q)) - h_{X,L}(\pi\circ\sigma^n(Q))\right| \\ &\leq 2\alpha^{-n}\kappa',\end{aligned}$$

since $\pi\circ\sigma^n(Q) = \pi\circ\sigma^n(P)$. In other words, given any $\varepsilon > 0$ there exists an n such that

$$\left|\hat{h}_{X,C,L}(P) - \hat{h}_{X,C,L}(Q)\right| < \varepsilon$$

for all $Q \in U_{P,n}$, whence the claim. □

Remark 3.4. In single-valued arithmetic dynamics [8], the canonical height is defined on the base variety X, where here we have defined it on the path space $\mathscr{P}$. In light of the above proposition, it is natural to define two real-valued functions on $X(\overline{K})$ itself by

$$\hat{h}_{X,C,L,\max}(a) = \max\left\{\hat{h}_{X,C,L}(P) : \pi(P) = a\right\}$$
$$\hat{h}_{X,C,L,\min}(a) = \min\left\{\hat{h}_{X,C,L}(P) : \pi(P) = a\right\}.$$

It is clear from Theorem 1.1 that both of these differ from any Weil height (on X relative to L) by a bounded amount, and many interesting questions about the canonical height can phrased in terms of $\hat{h}_{X,C,L,\min}$ in particular.

4. Specialization

We now state the specialization theorem, for which we must set the stage. Let K be a number field, let U/K be a curve, and let $k = \overline{K}$. By a *family of polarized correspondences* over U, we mean simply a variety $X \to U$ with a correspondence $C \subseteq X \times_U X$, a divisor $L \in \mathrm{Div}(X) \otimes \mathbb{R}$, and a real number $\alpha > 1$ such that $y^*L - \alpha x^*L$ is linearly equivalent to a divisor fibral to the projection $C \to U$. This is equivalent to requiring the generic fibre in the family to be a polarized correspondence. From this definition it follows easily that there is an affine open subset $U_0 \subseteq U$ such that for all $t \in U_0$ the triple of fibres (X_t, C_t, L_t) defines a polarized correspondence as defined above. We will always take $U_0 = U$, paring down to an open subset when necessary, although when we speak of heights on U we of course mean heights on a projective curve birational to U.

From Section 2, we can construct the path space $\mathscr{P}$ of the correspondence as a U-scheme. Just as in the arithmetic case, however, one should expect $\mathscr{P}(U)$ to

be rather small and uninteresting, and similarly for $\mathscr{P}(V)$ for any $V \to U$ of finite type.

In analogy to the arithmetic case, let $\varphi : \mathscr{U} \to U$ be a scheme not necessarily of finite type over U, with φ dominant, and consider paths $P \in \mathscr{P}(\mathscr{U})$ such that $\pi(P) \in X(U)$ (in other words, such that the map $\pi(P) : \mathscr{U} \to X$ factors through U). We can specialize such a path at points $t \in \mathscr{U}(k)$, and ask how $\hat{h}_{X_{\varphi(t)},C_{\varphi(t)},L_{\varphi(t)}}(P_t)$ varies as a function of t. We should note that the relative non-specificity of $\mathscr{U}$ is misleading. Let $\mathscr{U}' \to U$ be the limit of the inverse system of curves V with a finite map $V \to U$. Then it is straightforward to show that any $P : \mathscr{U} \to \mathscr{P}$ such that $\pi(P) : \mathscr{U} \to X$ factors through U, must in turn factor through some $P' : \mathscr{U}' \to \mathscr{P}$. Without loss of generality, then, we will always suppose that $\mathscr{U} = \mathscr{U}'$. This is the geometric equivalent of working over $\overline{K}$, rather than an arbitrary field extension of K.

Recall that the generic fibre (over U) of a family of correspondences as above is a polarized correspondence defined over the function field $k(U)$, which we denote by X_η, C_η, L_η. There is a natural height on $X_\eta(\overline{k(U)})$ relative to L_η, corresponding on $k(U)$ to the degree of a function on U, and it is easy to check that the construction of the canonical height in Section 3 works in this setting as well. In particular, given any finite cover $V \to U$ and any morphism $f : V \to X$, we have a corresponding point $f_\eta \in X_\eta(\overline{k(U)})$ given by some embedding $k(V) \to \overline{k(U)}$, and we define $h_{X_\eta,L_\eta}(f_\eta)$ accordingly; although the notation f_η is ambiguous, depending on the choice of embedding, the height is well-defined. We denote by $\hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta)$ the canonical height as computed on the generic fibre (noting that, since φ is dominant and η is the generic point of $\mathscr{U}$, $\varphi(\eta)$ is the generic point of U).

Theorem 4.1, generalizing a result of the same name due to Call and Silverman [8], establishes that the canonical height varies nicely in such a family.

Theorem 4.1. *In the setting above, let h_U be a Weil height on U corresponding to any divisor class of degree 1. Then*

$$\hat{h}_{X_{\varphi(t)},C_{\varphi(t)},L_{\varphi(t)}}(P_t) = \left(\hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta) + o(1)\right) h_U(\varphi(t)),$$

where $o(1) \to 0$ as $h_U(\varphi(t)) \to \infty$. In particular, if $\hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta) > 0$, then the set

$$\{\varphi(t) : P_t \text{ is finitely supported}\} \subseteq U(k)$$

is a set of bounded height.

The proof of the theorem comes in a series of lemmas, and follows the outline of the proof of Theorem 4.1 of [8]. The following lemma is the analogue of [8, Theorem 3.1], and is proven in the same way.

Lemma 4.2. *Let X, C, L be a family of polarized correspondences over U, and let h_U be a degree-one height on U. Then there exist constants c_1 and c_2 such that for all $t \in U(\overline{K})$ and all $P \in \mathscr{P}_t(\overline{K})$,*

$$\left|\hat{h}_{X_t,C_t,L_t}(P) - h_{X_t,L_t}(\pi(P))\right| \leq c_1 h_U(t) + c_2$$

Proof. Write $\tilde{U}$ for the projective closure of U. Note that this lemma is the analogue of [8, Theorem 3.1], but the proof is in fact simpler. We need not worry about resolving the indeterminacy of rational maps, since the coordinate projections are always morphisms. Now, by the hypothesis that $y^*L - \alpha x^*L$ is fibral, as in [8,

Proof of Theorem 3.1] there is an effective divisor D on $\tilde{U}$ such that for the map $u : C \to U$ we have

$$u^*D > y^*L - \alpha x^*L > -u^*D.$$

Choosing H on $\tilde{U}$ ample such that $nH + D$ and $nH - D$ are both ample, we have

$$|h_{C,y^*L-\alpha x^*L}(Q)| \leq n h_{\tilde{U},H}(u(Q)) + O(1)$$

for all Q on Y. Thus, for a given path $P \in \mathscr{P}(U)$,

$$\begin{aligned} |h_{X_t,L_t}(\pi \circ \sigma(P_t)) - \alpha h_{X_t,L_t}(\pi(P_t))| &= |h_{X,L}(y \circ \varepsilon(P_t)) - \alpha h_{X_t,L_t}(x \circ \varepsilon(P_t))| \\ &= |h_{Y,y^*L}(\varepsilon(P_t)) - \alpha h_{Y,x^*L}(\varepsilon(P_t))| \\ &= |h_{C,y^*L-\alpha x^*L}(\varepsilon(P_t))| \\ &\leq n h_{\tilde{U},H}(u \circ \varepsilon(P_t)) + O(1) \\ &= n h_{\tilde{U},H}(t) + O(1). \end{aligned}$$

The result now follows from the fact that $h_{\tilde{U},H}$ is bounded above by a multiple of any given degree one height on U. □

Lemma 4.3. *Let $\psi : V \to U$ be a finite cover, and let $f : V \to X$ be a morphism. Then for any $t \in V$,*

$$h_{X_{\psi(t)},L_{\psi(t)}}(f(t)) = h_{X_\eta,L_\eta}(f_\eta) h_U(\psi(t)) + O\left(\sqrt{h_U(\psi(t))}\right),$$

where the implied constant is allowed to depend on f.

Proof. Essentially by definition,

$$h_{X_\eta,L_\eta}(f_\eta) = \frac{\deg(f^*L)}{[k(V) : k(U)]} = \frac{\deg(f^*L)}{\deg(\psi)}.$$

Let D be the divisor relative to which h is defined, and let

$$E = \frac{1}{h_{X_\eta,L_\eta}(f_\eta)} f^*L,$$

so that $\deg(E) = \deg(D) = 1$. Note that, by [19, Proposition 5.4, p. 115], we have

$$h_{U,E}(t) = h_{U,D}(t) + O\left(\sqrt{h_{U,D}(t)}\right).$$

Also, the embedding $X_t \to X$ pulls L back to L_t, and so for any $t \in V$ we have,

$$\begin{aligned} h_{X_t,L_t}(f(t)) &= h_{X,L}(f(t)) \\ &= h_{V,f^*L}(t) + O(1) \\ &= h_{X_\eta,L_\eta}(f_\eta) h_{U,E}(t) + O(1) \\ &= h_{X_\eta,L_\eta}(f_\eta) h_{U,D}(\psi(t)) + O\left(\sqrt{h_{U,D}(\psi(t))}\right). \end{aligned}$$

□

Proof of Theorem 4.1. Let $P \in \mathscr{P}(\mathscr{U})$ such that $\pi(P) \in X(V)$, for some finite $\psi : V \to U$ that $\varphi : \mathscr{U} \to U$ factors as $\varphi = \psi \circ \xi$, for some $\xi : \mathscr{U} \to V$. We may specialize P at $t \in \mathscr{U}(\overline{K})$ to obtain a path $P_t \in \mathscr{P}_{\varphi(t)}(\overline{K})$. Note that, in this

setting, $\pi(P_t) = \pi(P)_{\xi(t)} \in X(V)$. Then interpreting $\pi(P)$ as a morphism from V to X, we have the following.

$$\begin{aligned}\Big|\hat{h}_{X_{\varphi(t)},C_{\varphi(t)},L_{\varphi(t)}}(P_t) - \hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta)h_U(\varphi(t))\Big| \\ \leq \Big|\hat{h}_{X_{\varphi(t)},C_{\varphi(t)},L_{\varphi(t)}}(P_t) - h_{X_{\varphi(t)},L_{\varphi(t)}}(\pi(P_t))\Big| \\ + \big|h_{X_{\varphi(t)},L_{\varphi(t)}}(\pi(P)_{\xi(t)}) - h_{X_{\varphi(\eta)},L_{\varphi(\eta)}}(\pi(P)_\eta)h_U(\psi\circ\xi(t))\big| \\ + \Big|h_{X_{\varphi(\eta)},L_{\varphi(\eta)}}(\pi(P)_\eta) - \hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta)\Big| h_U(\varphi(t)) \\ \leq c_3 h_U(\varphi(t)) + c_4\sqrt{h_U(\varphi(t))} + c_5,\end{aligned}$$

where c_4 and c_5 depend on V and P, but c_3 depends only on X, C, L, and U. We thus have

$$\limsup_{h_U(\varphi(t))\to\infty}\left|\frac{\hat{h}_{X_{\varphi(t)},C_{\varphi(t)},L_{\varphi(t)}}(P_t)}{h_U(\varphi(t))} - \hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta)\right| \leq c_3, \tag{4}$$

noting that this bound does not depend on P or the cover $V \to U$.

Now, given our path $P \in \mathscr{P}(\mathscr{U})$ satisfying $\pi(P) \in X(U)$, note that $Q = \sigma^n(P)$ is a path $Q \in \mathscr{P}(\mathscr{U})$ satisfying $\pi(Q) \in X(V)$ for some finite $\psi : V \to U$ as above. Thus we may apply the estimate (4) to the path $Q = \sigma^n(P)$, to obtain

$$\begin{aligned}&\limsup_{h(\varphi(t))\to\infty}\left|\frac{\hat{h}_{X_{\varphi(t)},C_{\varphi(t)},L_{\varphi(t)}}(P_t)}{h_U(\varphi(t))} - \hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta)\right| \\ &\quad= \limsup_{h_U(\varphi(t))\to\infty}\alpha^{-n}\left|\frac{\hat{h}_{X_{\varphi(t)},C_{\varphi(t)},L_{\varphi(t)}}(Q_t)}{h_U(\varphi(t))} - \hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(Q_\eta)\right| \leq \alpha^{-n}c_3.\end{aligned}$$

Since this holds for all n, and since the absolute value is non-negative, we have in fact

$$\lim_{h_U(\varphi(t))\to\infty}\left|\frac{\hat{h}_{X_{\varphi(t)},C_{\varphi(t)},L_{\varphi(t)}}(P_t)}{h_U(\varphi(t))} - \hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta)\right| = 0.$$

This proves the main claim.

If the set

$$\{\varphi(t) : P_t \text{ is finitely supported}\} \subseteq U(k)$$

is not of bounded height, then we may choose a sequence $t_n \in \mathscr{U}(k)$ such that P_{t_n} is finitely supported for all n, and $h_U(\varphi(t_n)) \to \infty$ with n. But then we have

$$0 = \lim_{n\to\infty}\left|\frac{\hat{h}_{X_{\varphi(t)},C_{\varphi(t)},L_{\varphi(t)}}(P_t)}{h_U(\varphi(t))} - \hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta)\right| = \hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta)$$

since the fractional term always has vanishing numerator. □

As a corollary to the theorem, we note the following.

Corollary 4.4. *In the setting of Theorem 4.1, consider a dominant morphism $\varphi : \mathscr{U} \to U$ and a path $P \in \mathscr{P}(\mathscr{U})$ with $\hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta) > 0$. Then the set*

$$\{t \in U(K) : P_s \text{ is finitely supported for some } \varphi(s) = t\}$$

is a finite set.

It would be of significant interest to determine the conditions under which $\hat{h}_{X_{\varphi(\eta)},C_{\varphi(\eta)},L_{\varphi(\eta)}}(P_\eta) = 0$, but experience indicates that this is likely more challenging than the yet-unresolved number field case [4, 6].

5. Measured sets of absolute values and augmented divisors

In this section we lay out some of the basics needed for our discussion of local canonical heights. The fundamental difference between this treatment and that in [13, 19] is that we work in a fundamental way over the algebraic closure $k = \overline{K}$ of a number field, which bars us from expressing global heights as weighted sums of local heights in the usual way. The remedy is to employ Gubler's theory of M-fields [12] although, since the author developed this material independently before learning of Gubler's work, the presentation here differs slightly from [12].

Let k be a field. A *measured set of absolute values* on k will be a set M_k of absolute values, along with a σ-algebra $\mathcal{B}$ of subsets, and a measure μ, such that for every non-zero $\alpha \in k$, the function $v \mapsto \log |\alpha|_v$ is measurable. We will say that such an object *satisfies the product formula* if and only if we have

$$\int_{M_k} \log |\alpha|_v d\mu(v) = 0$$

for any $\alpha \in k^*$.

Here is a simple example to aid the reader's intuition.

Example 5.1. Let $k = \mathbb{Q}$, let $M_{\mathbb{Q}}$ be the usual set of absolute values, let $\mathcal{B}$ be the set of all subsets of $M_{\mathbb{Q}}$, and let μ be the counting measure. Then $v \mapsto \log |\alpha|_v$ is measurable for any $\alpha \in \mathbb{Q}^*$, since the function vanishes off of a finite set, and

$$\int_{M_k} \log |\alpha|_v d\mu(v) = \sum_{v \in M_{\mathbb{Q}}} \log |\alpha|_v = 0$$

by the standard product formula.

More generally, if k is a number field, M_k is the usual set of absolute values, $\mathcal{B}$ is the power set of M_k, and μ is a weighted counting measure assigning weight $[k_v : \mathbb{Q}_v]/[k : \mathbb{Q}]$ to each $v \in M_k$, then this is a measured set of absolute values on k satisfying the product formula.

We begin with a theorem ensuring that, for any number field K, there is a natural way to view the set M_k of places of $k = \overline{K}$ as a measured set of absolute values. In order to make the result somewhat more generally applicable, however, we relax the condition that K is a number field. We say that the set of absolute values M_K on the field K is *coherent* if and only if, for every finite extension $E/L/K$ and every $v \in M_L$, we have

$$\sum_{w|v} [E_w : L_w] = [E : L], \tag{5}$$

where the sum is over extensions w of v to E. For instance, the usual set of absolute values on a number field is coherent, as is the usual set of absolute values on a function field in characteristic zero.

Theorem 5.2. *Let K be a field, let M_K be a coherent set of absolute values on K, and let M_k be the set of absolute values on $k = \overline{K}$ which restrict on K to elements*

of M_K. Then we may equip M_k with the structure of a measured set of absolute values in such a way that for any finite L/K and $w \in M_k$, the set

$$A_{w,L} = \{v \in M_k : v = w \text{ on } L\}$$

is measurable, and

$$\mu(A_{w,L}) = \frac{[L_w : K_w]}{[L : K]}.$$

Proof. Fix an absolute value $v \in M_K$, and let $M_{k,v}$ be the set of absolute values on k restricting on K to v. We write $\mathscr{F}_{0,v}$ for the algebra of sets generated by sets of the form $A_{w,L}$, for L/K finite and $w \in M_L$, and define μ_v on $\mathscr{F}_{0,v}$ as in the statement of the theorem, extended to disjoint unions by additivity. Note that the fact that μ_v is well-defined is precisely the coherence condition imposed on the places of K. In particular, if $E/L/K$ is a tower of finite extensions, then

$$\mu_v(A_{w,L}) = \frac{[L_w : K_w]}{[L : K]} = \sum_{\substack{u \in M_E \\ u = w \text{ on } L}} \frac{[E_u : K_u]}{[E : K]} = \sum_{\substack{u \in M_E \\ u = w \text{ on } L}} \mu_v(A_{u,L}).$$

We wish to show that μ_v extends to a measure on some σ-algebra, and so we recall a special case of a theorem of Carathéodory.

Theorem 5.3 (Carathéodory Extension Theorem [**?**, p. 97])**.** *Let X be a set, let $\mathcal{B}$ be an algebra of subsets of X, and let $\mu : \mathcal{B} \to \mathbb{R}$ be a function such that $\mu(\emptyset) = 0$, and*

$$\mu\left(\bigcup_{i \geq 0} B_i\right) = \sum_{i \geq 0} \mu(B_i) \tag{6}$$

for any sequence $B_i \in \mathcal{B}$ of disjoint sets whose union happens to be in $\mathcal{B}$. Then μ extends to a measure on a σ-algebra containing $\mathcal{B}$. Furthermore, if μ is σ-finite on $\mathcal{B}$, then this extension is unique.

Remark 5.4. Carathéodory's theorem is proven by extending μ to an outer measure μ^* on X. Taking $\mathscr{F}$ to be the collection of subsets $A \subseteq X$ satisfying

$$\mu^*(E) = \mu^*(E \cap A) + \mu^*(E \setminus A)$$

for every $E \subseteq X$, it is shown that $\mathcal{B} \subseteq \mathscr{F}$, and that $\mathscr{F}$ is the maximal σ-algebra to which μ^* restricts as a measure. If we take $\mathscr{F}$ to be this maximal σ-algebra (rather than some smaller σ-algebra containing $\mathcal{B}$), then the construction is unique.

Lemma 5.5. *The function μ_v extends uniquely to a measure on some σ-algebra $\mathscr{F}_v$ containing $\mathscr{F}_{0,v}$.*

Proof. We need only show that $\mathscr{F}_{0,v}$ and μ meet the conditions of Theorem 5.3, which really comes down to establishing (6). The facts that $\mu_v(\emptyset) = 0$ and $\mu_v(M_{k,v}) = 1$ are immediate, while the (finite) additivity of μ_v follows immediately from (5). The σ-additivity is the main content of the lemma.

Note that, if we take the sets $A_{w,L}$ as the basis for a topology on $M_{k,v}$, then this space is naturally the inverse limit of the spaces $M_{L,v}$ with the discrete topology, for L/K finite. Since these spaces are compact and Hausdorff, so is $M_{k,v}$. Now, suppose we have $B_i \in \mathscr{F}_{0,v}$. Note that the elements of $\mathscr{F}_{0,v}$ are both open and closed, and so if $\bigcup_{i \geq 0} B_i \in \mathscr{F}_{0,v}$, there must be a finite subcollection of the B_i with

the same union. But these sets are disjoint, and hence all but finitely many are empty. Now σ-additivity follows from finite additivity. □

We now show quickly that one may construct a disjoint union of a collection of measure spaces.

Lemma 5.6. *Let X_i be disjoint sets, for $i \in I$, let $\mathscr{F}_i$ be a σ-algebra on X_i, and let μ_i be a measure on $\mathscr{F}_i$. Then there is a σ-algebra $\mathscr{F}$ on $X = \bigcup X_i$ and a measure μ on $\mathscr{F}$ such that, for each i, $\mathscr{F}_i \subseteq \mathscr{F}$ and μ_i is the restriction to $\mathscr{F}_i$ of μ.*

Proof. Let $\mathscr{F}$ be the collection of subsets $B \subseteq X$ with the property that $B \cap X_i \in \mathscr{F}_i$ for all i. It is easy to check that $\mathscr{F}$ is a σ-algebra, and of course $\mathscr{F}_i \subseteq \mathscr{F}$ for all i. We can define μ on $\mathscr{F}$ by

$$\mu(B) = \sum_{i \in I} \mu_i(B \cap X_i).$$

Checking that this is a measure involves interchanging two sums, both of non-negative real numbers. □

Applying the construction given by Lemma 5.6 to produce from the probability spaces $(M_{k,v}, \mathscr{F}_v, \mu_v)$ a measure space $(M_k, \mathscr{F}, \mu)$, we have proven the theorem. Note that, if M_K is countable (as is the case when K is a number field, or the function field of a curve over a countable field), then this measure space is a countable disjoint union of probability spaces (the spaces $(M_{k,v}, \mathscr{F}_v, \mu_v)$ each have total mass 1), and hence is σ-finite. □

Given the notion of a measured set of absolute values, and the existence of a natural such structure on $\overline{K}$ for any number field K, we formulate an appropriate notion of local heights on certain sorts of schemes. Our notion is motivated by the construction via Néron divisors [19, Chapter 10], but we introduce a class of objects which do not quite restrict to the notion of a Néron divisor on a variety. It would be possible to complete much of this construction with a more direct generalization of Néron divisors, but the measure-theoretic context makes another construction more natural.

Given a measured set M_k of absolute values on k, the appropriate analogue of an M_K-constant in the sense of [19] is an element of $L^1(M_k)$. Note that such functions are not precisely M_K-constants, even on $M_{\mathbb{Q}}$ with the counting measure. In particular, if we enumerate the places of $\mathbb{Q}$ as v_1 the archimedean place, v_2 the 2-adic place, v_3 the 3-adic place, *et cetera*, then $\gamma(v_n) = 2^{-n} \in L^1(M_{\mathbb{Q}})$. But γ is not finitely supported, and hence is not an $M_{\mathbb{Q}}$-constant in the sense of [19]. On the other hand, if K is a number field and γ is an M_K-constant in the sense of [19], then $\gamma \in L^1(M_K)$ simply because γ is finitely supported.

As usual, a *Cartier divisor* on a k-scheme X is represented by a cover U_i, for $i \in I$, by affine open sets, along with a collection of rational functions $f_i \in k(U_i)$, subject to the condition that f_i/f_j and f_j/f_i are both regular on $U_i \cap U_j$, for any choice of indices. As is conventional, we take a Cartier divisor to be a maximal such data, with the Cartier divisors represented by (U_i, f_i) and (U_i, g_i) identified if and only if f_i/g_i and g_i/f_i are regular on each U_i. The zero divisor will then be defined by $f_U = 1$ for all $U \subseteq X$. The *support* $\mathrm{Supp}(D)$ is defined by the condition that $P \in U \subseteq X$ is in the support of D if and only if f_U or $1/f_U$ is not regular at P.

We now define an *augmented divisor* to be a Cartier divisor (U_i, f_i) with an additional piece of information $\alpha_i : U_i(k) \times M_k \to \mathbb{R}$ for each $i \in I$ such that

(1) $\alpha_i(\cdot, v) : U_i(k) \to \mathbb{R}$ is v-adically continuous for each $v \in M_k$,
(2) $\alpha_i(P, \cdot)$ is measurable for each $P \in U_i(k)$,
(3) $|\alpha_i(P, v)| \leq \gamma(v)$ for some $\gamma \in L^1(M_k)$ and all $P \in U_i(k), v \in M_k$,
(4) $\log |f_i(P)/f_j(P)|_v = \alpha_i(P, v) - \alpha_j(P, v)$ on $U_i \cap U_j$.

Note that, if k is a number field equipped with the above weighted counting measure, then every Néron divisor in the sense of [19, Chapter 10] is an augmented divisor. We shall abuse notation and use the same symbol for an augmented divisor as for its underlying Cartier divisor.

To each augmented divisor $D = (U_i, f_i, \alpha_i)$, we associate a *local height* (i.e., a Weil function)

$$\lambda_{X,D} : (X(k) \setminus \mathrm{Supp}(D)) \times M_k \to \mathbb{R}$$

by

$$\lambda_{X,D}(P, v) = -\log |f_i(P)|_v + \alpha_i(P, v)$$

for any $P \in U_i$. That this is well-defined follows from item (4) in the definition.

Given a scheme X over an algebraically closed field k with a measured set of absolute values, we will write $\mathrm{Div}(X)$ for the set of Cartier divisors, and $\mathrm{ADiv}(X)$ for the set of augmented divisors. It is clear that both are abelian groups in a natural way, and that there is a homomorphism

$$\mathrm{ADiv}(X) \to \mathrm{Div}(X) \tag{7}$$

obtained by forgetting the functions α_i. As remarked, in the case where X is a projective variety we have $\mathrm{NDiv}(X) \subseteq \mathrm{ADiv}(X)$ in a natural way, where $\mathrm{NDiv}(X)$ denotes the group of Néron divisors, and the fact that $\mathrm{NDiv}(X) \to \mathrm{Div}(X)$ is surjective is more-or-less the content of the local height machine. It is then of great interest to determine when the corresponding map (7) is surjective.

In order to give an answer to this question in the relevant case, we will say that a k-scheme $\mathscr{X}$ is a *provariety* if and only if it can be exhibited as the limit of an inverse system of projective varieties over k with finite transition maps.

Lemma 5.7. *Let $\mathscr{X}$ be a provariety over k. Then every Cartier divisor on $\mathscr{X}$ is the pull-back of a Cartier divisor through some morphism $f : \mathscr{X} \to X$ to a projective variety. As a consequence, the homomorphism $\mathrm{ADiv}(\mathscr{X}) \to \mathrm{Div}(\mathscr{X})$ is surjective.*

Proof. Suppose that we are given a Cartier divisor $\mathscr{D}$ on $\mathscr{X}$ by the data (U_t, g_t), with indices $t \in T$. Let I be a partially ordered set, and let (X_i, f_{ij}) be an inverse system of varieties over k such that $\mathscr{X} = \varprojlim X_i$. We have canonical (affine) maps $f_i : \mathscr{X} \to X_i$, for each $i \in I$. By [18, Chapter 31, Lemma 3.8] each U_t can be expressed as $U_t = f_{i_t}^{-1}(V_{i_t})$ for some $i_t \in I$ and some affine open $V_{i_t} \subseteq X_{i_t}$. If we pare down the U_t to a finite cover, all of these maps factor through some particular index $0 \in I$, and so we may write $U_t = f_0^{-1}(V_t)$ for some affine opens $V_t \subseteq X_0$.

Now, since the sheaf of rational functions on $\mathscr{X}$ is the colimit of the sheaves of rational functions on the X_i (essentially by definition [18, Chapter 31, Lemma 2.2]), each $g_t \in k(U_t)$ must be the pull-back to U_t of some rational function $h_{i_t} \in k(f_{i_t 0}^{-1}(U_t))$. Again, now that the index set T is finite, we can assume without loss of generality that $i_t = 0$ for all t, and so $g_t \in k(U_t)$ is the pull-back of some

$h_t \in k(V_t)$. But now we have some data (V_t, h_t) representing a Cartier divisor D on some X_0, and by construction $\mathscr{D} = f_0^*(D)$.

Since X_0 is a projective variety, the Cartier divisor D on X_0 admits the structure of a Néron divisor [19, Theorem 3.5, p. 261]. We may simply pull back this additional structure to $\mathscr{D}$. If $\beta_t : V_t \times M_k \to \mathbb{R}$ are the functions making D a Néron divisor, it is easy to check that $\alpha_t(P, v) = \beta_t(f(P), v)$ gives a collection of functions α_t giving $\mathscr{D}$ the structure of an augmented divisor. □

Remark 5.8. For the purposes of this section, it would be sufficient to define a Néron divisor on $\mathscr{X}$ to be one that is the pull-back of a Néron divisor on some projective variety $\mathscr{X} \to X$. This approach seems less natural given the measure-theoretic setting, but also causes significant problems in the application of Lemma 6.2 below.

Lemma 5.9. *Let $\mathscr{X}$ be a provariety over k. If $\mathscr{D}$ is an augmented divisor whose corresponding Cartier divisor is trivial, then $\lambda_{\mathscr{X},\mathscr{D}}(P, \cdot) \in L^1(M_k)$ for each $P \in \mathscr{X} \setminus \mathrm{Supp}(D)$, and the value of the integral*

$$\int_{M_k} \lambda_{\mathscr{X},\mathscr{D}}(P, v) d\mu(v)$$

is bounded independent of P.

Proof. Let U be an affine open on which $\mathscr{D}$ is given by the data (U, f, α). Note that since $\mathscr{D}$ is trivial, both f and $1/f$ are regular on U. As in the proof of the previous lemma, we may choose a projective variety X, an affine open V, a rational function $h \in k(V)$, and a morphism $g : \mathscr{X} \to X$ such that $U = g^{-1}(V)$, and such that $f = h \circ g$. Now, $\log |f(P)|_v = \log |h(g(P))|_v$, and so if X, V, h, and $g(P)$ are defined over the finite extension L/K then this function is constant on sets of the form $A_{w,L}$. Since it also vanishes off of a finite collection of sets of this form, it is L^1. The function $\alpha(P, \cdot)$ is measurable and bounded by an L^1 function, and hence is L^1. Now $\lambda_{X,D}(P, \cdot)$ is the sum of two L^1 functions, and hence is itself L^1. Similarly, the fact that $\lambda_{X,D}(\cdot, v)$ is v-adically continuous on U follows from the same fact for $\log |f(\cdot)|_v$ and $\alpha(\cdot, v)$.

Now, since f is regular at P, the product formula gives

$$\int_{M_k} \log |f(P)|_v d\mu(v) = 0.$$

It follows that,

$$\begin{aligned} \left| \int_{M_k} \lambda_{\mathscr{X},D}(P, \cdot) d\mu(v) \right| &\leq \left| \int_{M_k} -\log |f(P)|_v d\mu(v) \right| + \left| \int_{M_k} \alpha(P, v) d\mu(v) \right| \\ &\leq \int_{M_k} |\alpha(P, v)| \, d\mu(v) \\ &\leq \int_{M_k} \gamma(v) d\mu(v), \end{aligned}$$

where γ is the L^1 function ensured by the definition. This bound is independent of P, and hence of U. □

In light of Lemma 5.7, for any Cartier divisor $\mathscr{D}$ on $\mathscr{X}$, we may define

$$h_{\mathscr{X},\mathscr{D}}(P) = h_{X,D}(f(P))$$

for any projective variety $f : \mathscr{X} \to X$ and Cartier divisor D such that $\mathscr{D} = f^*D$. The proof of Lemma 5.7 ensures that any divisor on $\mathscr{X}$ can be so exhibited, and the usual functoriality properties of the Weil height [13, Theorem B.3.2, p. 184] show that the definition is independent of the choice of X, at least up to a bounded error term. Indeed, all of the basic properties one would like about heights on provarieties follow quickly from the corresponding properties for varieties.

Remark 5.10. The observation that heights on provarieties have essentially the same properties as heights on projective varieties gives another approach to the construction of the canonical height $\hat{h}_{X,C,L}$. Note that the relation $y^*L \sim \alpha x^*L$, in concert with the commutative diagram in Theorem 2.1 implies

$$\sigma^*\mathscr{L} = \sigma^*\pi^*L = \varepsilon^*y^*L \sim \varepsilon^*\left(\alpha x^*L\right) = \alpha\pi^*L = \alpha\mathscr{L}.$$

In particular, one has a morphism of k-schemes $\sigma : \mathscr{P} \to \mathscr{P}$ satisfying

$$h_{\mathscr{P},\mathscr{L}}(\sigma(P)) = \alpha h_{\mathscr{P},\mathscr{L}}(P) + O(1),$$

and one can dispatch with the correspondence entirely, and use the (suddenly applicable) arguments in [8] to construct the canonical height.

We have avoided this approach both because it takes the emphasis off of the underlying correspondence, and also because the machinery required by the arguments in Section 3 is far more elementary.

Theorem 5.11. *Let $\mathscr{X}$ be a provariety over k, and let $\mathscr{D}$ be an augmented divisor on $\mathscr{X}$. Then we have*

$$h_{\mathscr{X},\mathscr{D}}(P) = \int_{M_k} \lambda_{\mathscr{X},\mathscr{D}}(P,v)d\mu(v) + O(1),$$

for all $P \notin \mathrm{Supp}(\mathscr{D})$.

Proof. As shown in the proof of Lemma 5.7, there exists a projective variety $f : \mathscr{X} \to X$ and a divisor D on X such that $\mathscr{D} = f^*D$. If $\lambda_{X,D}$ is the local height associated to some Néron divisor on D, and if X, $f(P)$, and D are defined over the finite extension L/K, then we have

$$\begin{aligned}\int_{M_k} \lambda_{X,D}(f(P),v)d\mu(v) &= \sum_{v\in M_L} \frac{[L_v : K_v]}{[L:K]}\lambda_{X,D}(f(P),v) \\ &= h_{X,D}(f(P)),\end{aligned}$$

since the function $\lambda_{X,D}(f(P),\cdot)$ is constant on sets of the form $A_{v,L}$ for $v \in M_L$, in the notation of Theorem 5.2. On the other hand, we have $h_{\mathscr{X},\mathscr{D}}(P) = h_{X,D}(f(P)) + O(1)$ by definition. It thus suffices to show that

$$\int_{M_k} \left(\lambda_{\mathscr{X},\mathscr{D}}(P,v) - \lambda_{X,D}(f(P),v)\right) d\mu(v)$$

is bounded independent of P. But this follows from Lemma 5.9. Specifically, $(P,v) \mapsto \lambda_{X,D}(f(P),\cdot)$ is a local height on $\mathscr{X}$ relative to $\mathscr{D} = f^*D$, and so the difference $(P,v) \mapsto \lambda_{\mathscr{X},\mathscr{D}}(P,v) - \lambda_{X,D}(f(P),v)$ is a local height on $\mathscr{X}$ relative to some augmented divisor $\mathscr{E}$ whose underlying Cartier divisor is trivial (simply because the forgetful map $\mathrm{ADiv}(X) \to \mathrm{Div}(X)$ is a homomorphism). □

6. Local canonical heights

In this section we construct local canonical heights for a polarized correspondence defined over the algebraic closure k of a number field. We equip M_k with the structure of a σ-finite measure space, with measure μ, as in the previous section. In order to state the theorem, let X be a projective variety over k, and let C be a correspondence on X polarized by L and $\alpha > 1$. Construct $\mathscr{P}$ as in Section 2, and to simplify notation let $\mathscr{L} = \pi^* L$.

Theorem 6.1. *Given the data above, there exists a function*

$$\hat{\lambda}_{X,C,L} : (\mathscr{P}(k) \setminus \mathrm{Supp}(\mathscr{L})) \times M_k \to \mathbb{R}$$

such that the following hold:

1. *The function $\hat{\lambda}_{X,C,L}$ is the local height function associated to some augmented divisor on $\mathscr{P}$ with Cartier divisor $\mathscr{L}$.*
2. *For some rational function $f \in k(C)^* \otimes \mathbb{R}$, for all $P \in \mathscr{P}(k)$ with $\varepsilon(P)$ not a zero or pole of f, and all $v \in M_k$,*
$$\hat{\lambda}_{X,C,L}(\sigma(P), v) = \alpha \hat{\lambda}_{X,C,L}(P, v) - \log |f \circ \varepsilon(P)|_v.$$
3. *For all $P \in \mathscr{P}(k)$,*
$$\hat{h}_{X,C,L}(P) = \int_{M_k} \hat{\lambda}_{X,C,L}(P, v) d\mu(v).$$

The proof requires a lemma extending an argument dating back to Tate, present also in [8].

Lemma 6.2. *Fix a set $\mathcal{X}$, a measure space $\mathcal{Y}$, and a real number $\alpha > 1$. Suppose also that we are given a topology on each fibre of $\mathcal{X} \times \mathcal{Y}$, and a function $\varphi : \mathcal{X} \to \mathcal{X}$ such that the induced function from $\mathcal{X} \times \mathcal{Y}$ to itself is continuous on each fibre. Finally, suppose we are given a function $\gamma : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ which*

1. *is bounded;*
2. *is continuous on fibres of $\mathcal{X} \times \mathcal{Y} \to \mathcal{Y}$;*
3. *is measurable on fibres of $\mathcal{X} \times \mathcal{Y} \to \mathcal{X}$; and*
4. *vanishes off of $\mathcal{X} \times W$, for some set $W \subseteq \mathcal{Y}$ of finite measure.*

Then there exists a unique bounded function $\hat{\gamma} : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, continuous in the first coordinate and integrable in the second, with

$$\gamma(x, y) = \hat{\gamma}(\varphi(x), y) - \alpha \hat{\gamma}(x, y) \tag{8}$$

for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$.

Proof. The construction is based on [8, Lemma 2.2]. First, let $\mathcal{Z}$ be the set of functions $F : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ which are bounded, continuous on fibres of $\mathcal{X} \times \mathcal{Y} \to \mathcal{Y}$, and measurable fibres of $\mathcal{X} \times \mathcal{Y} \to \mathcal{X}$. Note that $\gamma \in \mathcal{Z}$ by hypothesis, and the supremum norm gives $\mathcal{Z}$ the structure of a Banach space.

As in [8, Proof of Lemma 2.2], we let $\hat{\gamma} \in \mathcal{Z}$ denote the unique fixed point of the contraction $\delta \mapsto S(\delta)$ on $\mathcal{Z}$ given by

$$S(\delta)(x, y) = \frac{1}{\alpha} \left(\delta(\varphi(x), y) - \gamma(x, y) \right).$$

The relation (8) holds by definition. It remains to show that $\hat{\gamma}$ is in fact integrable (not merely measurable) in the second coordinate.

Choose $W \subseteq \mathcal{Y}$ of finite measure such that γ vanishes off of $\mathcal{X} \times W$. Then the same is true for $S^n(\gamma)$, for every n, and hence for $\hat{\gamma}$. Any function which is bounded, measurable, and supported on a set of finite measure is integrable. □

Proof of Theorem 1.8. Let $\lambda_{X,L}$ be some local height function, given by a Néron divisor, and let $f \in k(C)^* \otimes \mathbb{R}$ be a rational function with

$$\operatorname{div}(f) = y^*L - \alpha x^*L.$$

Note that, by the basic properties of local heights [13, Theorem B.8.1, p. 239] we have

$$\begin{aligned}
\lambda_{X,L}(\pi \circ \sigma(P), v) &= \lambda_{X,L}(y \circ \varepsilon(P), v) \\
&= \lambda_{C,y^*L}(\varepsilon(P), v) + O(1) \\
&= \lambda_{C,\alpha x^*L + \operatorname{div}(f)}(\varepsilon(P), v) + O(1) \\
&= \lambda_{C,\alpha x^*L}(\varepsilon(P), v) + \lambda_{C,\operatorname{div}(f)}(\varepsilon(P), v) + O(1) \\
&= \alpha\lambda_{X,L}(x \circ \varepsilon(P), v) + \lambda_{C,\operatorname{div}(f)}(\varepsilon(P), v) + O(1) \\
&= \alpha\lambda_{X,L}(\pi(P), v) - \log|f \circ \varepsilon(P)|_v + O(1),
\end{aligned}$$

for any $P \in \mathscr{P}(k)$ so long as $\pi(P), \pi \circ \sigma(P)$ are not in the support of L, and f and $1/f$ are regular at $\varepsilon(P)$. Note that the bound implied by the $O(1)$ notation is independent of P. Also, since we are applying [13, Theorem B.8.1, p. 239] in some number field L over which X, C, L, f are all defined, we may take the $O(1)$ bound to be constant on sets of the form $A_{w,L} \subseteq M_k$, and to vanish on all but finitely many of those. In other words, if we set

$$\gamma(P, v) = \lambda_{X,L}(\pi \circ \sigma(P), v) - \alpha\lambda_{X,L}(\pi(P), v) + \log|f \circ \varepsilon(P)|_v,$$

then we have just shown γ to be a bounded, measurable function which vanishes off $(\mathscr{P}(k) \setminus \operatorname{Supp}(\mathscr{L}) \cup \operatorname{Supp}(\sigma^*\mathscr{L})) \times W$, for some set W of finite measure. We wish to show that γ extends to a function on $\mathscr{P}(k) \times M_k$ which meets the conditions of Lemma 6.2.

Let $U \subseteq X$ be an affine open on which L is defined locally by $g = 0$, and let $\lambda_{X,L}$ be defined by

$$\lambda_{X,L}(P, v) = -\log|g(P)|_v + \beta(P, v).$$

Then, where defined,

$$\begin{aligned}
\gamma(P, v) &= -\log|g \circ \pi \circ \sigma(P)|_v + \beta(\pi \circ \sigma(P), v) + \alpha \log|g \circ \pi(P)|_v \\
&\quad -\alpha\beta(\pi(P), v) - \log|f \circ \varepsilon(P)|_v \\
&= \log|h(P)|_v + \beta(\pi \circ \sigma(P), v) - \alpha\beta(\pi(P), v) \qquad (9)
\end{aligned}$$

for

$$h = \frac{(g \circ \pi)^\alpha}{(f \circ \varepsilon)\,(g \circ \pi \circ \sigma)}.$$

But by hypothesis, h and $1/h$ are regular on U (since this function defines the trivial Cartier divisor), and so (9) gives an extension of γ to a bounded function on all of $U \times M_k$, continuous in the first variable and measurable in the second. Covering $\mathscr{P}$ by finitely many affine opens allows us to so extend γ to all of $\mathscr{P}(k) \times M_k$.

We may now apply Lemma 6.2 to γ to produce a unique bounded function $\hat{\gamma} : \mathscr{P}(k) \times M_k \to \mathbb{R}$, continuous in the first variable and integrable in the second, such that

$$\gamma(P, v) = \hat{\gamma}(\sigma(P), v) - \alpha\hat{\gamma}(P, v).$$

Given this, we set

$$\hat{\lambda}_{X,C,L}(P,v) = \lambda_{X,L}(\pi(P),v) - \hat{\gamma}(P,v).$$

Immediately from the definition we have

$$\begin{aligned}\hat{\lambda}_{X,C,L}(\sigma(P),v) &= \lambda_{X,L}(\pi\circ\sigma(P),v) - \hat{\gamma}(\sigma(P),v)\\ &= \alpha\lambda_{X,L}(\pi(P),v) - \log|f\circ\varepsilon(P))|_v + \gamma(P,v) - \hat{\gamma}(\sigma(P),v)\\ &= \alpha\lambda_{X,L}(\pi(P),v) - \log|f\circ\varepsilon(P))|_v - \alpha\hat{\gamma}(P,v)\\ &= \alpha\hat{\lambda}_{x,C,L}(P,v) - \log|f\circ\varepsilon(P))|_v\end{aligned}$$

wherever both sides are defined. We also note that, since $\hat{\gamma}$ is continuous in the first coordinate, measurable in the second, and bounded by an L^1 function of the second coordinate, the function $\hat{\lambda}_{X,C,L}$ is the local height associated to some augmented divisor with underlying Cartier divisor L.

We now aim to prove that these local canonical heights indeed sum to the global canonical height, following the same argument as in [8, Theorem 2.3]. We may define

$$h^*(P) = \int_{M_k} \hat{\lambda}_{X,C,L}(P,v)d\mu(v) \tag{10}$$

for $P \notin \mathrm{Supp}(\mathscr{L})$. We wish to show that this function extends to $\mathscr{P}(k)$.

Note that, if $g \in k(X)$, then $L' = L - \mathrm{div}(g)$ also offers a polarization of the correspondence in question. If we take as a local height $\lambda_{X,L-\mathrm{div}(g)}$ the function

$$\lambda_{X,L'}(P,v) = \lambda_{X,L}(P,v) + \log|g(P)|_v,$$

then we see that

$$y^*L' - \alpha x^*L' = y^*L - \alpha x^*L - y^*\,\mathrm{div}(g) + \alpha x^*\,\mathrm{div}(g) = \mathrm{div}(f'),$$

for $f' = (g\circ x)^\alpha f/g\circ y$. From this, note that defining γ as above relative to this divisor gives

$$\begin{aligned}\gamma'(P,v) &= \lambda_{X,L'}(\pi\circ\sigma(P),v) - \alpha\lambda_{X,L'}(\pi(P),v) + \log|f'\circ\varepsilon(P)|_v\\ &= \lambda_{X,L}(\pi\circ\sigma(P),v) + \log|g\circ\pi\circ\sigma(P)|v - \alpha\lambda_{X,L}(\pi(P),v)\\ &\quad -\alpha\log|g\circ\pi(P)|_v + \log\left|\frac{(f\circ\varepsilon(P))(g\circ x\circ\varepsilon(P))^\alpha}{(g\circ y\circ\varepsilon(P))}\right|_v\\ &= \lambda_{X,L}(\pi\circ\sigma(P),v) - \alpha\lambda_{X,L}(\pi(P),v) + \log|f\circ\varepsilon(P)|_v\\ &= \gamma(P,v),\end{aligned}$$

Since $\pi\circ\sigma = y\circ\varepsilon$ and $\pi = x\circ\varepsilon$. In other words, the function γ constructed above depends only on the linear equivalence class of L, and hence so does the function $\hat{\gamma}$. This gives

$$\hat{\lambda}_{X,C,L'}(P,v) = \hat{\lambda}_{X,C,L}(P,v) + \log|g(P)|_v,$$

wherever both sides are defined, and hence by the product formula

$$\int_{M_k} \hat{\lambda}_{X,C,L}(P,v)d\mu(v) = \int_{M_k} \hat{\lambda}_{X,C,L'}(P,v)d\mu(v)$$

(for any $P \notin \mathrm{Supp}(\mathscr{L}) \cup \mathrm{Supp}(\mathscr{L}')$).

By a standard argument, for any given $P \in \mathscr{P}(k)$ we can find some $L' \sim L$ such that $\pi(P) \notin \mathrm{Supp}(L')$, and hence (10) extends to a well-defined function

$h^* : \mathscr{P}(k) \to \mathbb{R}$. Now, since $\hat{\lambda}_{X,C,L}$ is in particular a local height on $\mathscr{P}$ relative to $\mathscr{L}$, we have by Theorem 5.11 and Theorem 1.1 that

$$h^*(P) = \int_{M_k} \hat{\lambda}_{X,C,L}(P,v)d\mu(v) = h_{\mathscr{P},\mathscr{L}}(P) + O(1) = \hat{h}_{X,C,L}(P) + O(1). \tag{11}$$

But note that, for

$$P \notin \mathrm{Supp}(\mathscr{L}) \cup \mathrm{Supp}(\sigma^*\mathscr{L}) \tag{12}$$

we have

$$\begin{aligned} h^*(\sigma(P)) &= \int_{M_k} \hat{\lambda}_{X,C,L}(\sigma(P),v)d\mu(v) \\ &= \int_{M_k} \alpha\hat{\lambda}_{X,C,L}(P,v)d\mu(v) - \int_{M_k} \log|f_i \circ \varepsilon(P)|_v d\mu(v) \\ &= \alpha h^*(P). \end{aligned}$$

As usual, we can always replace $\mathscr{L}$ by a linearly equivalent divisor so that all terms are defined at P, and hence we have $h^*(\sigma(P)) = \alpha h^*(P)$ for all $P \in \mathscr{P}(k)$. This in turn yields

$$\left| h^*(P) - \hat{h}_{X,C,L}(P) \right| = \alpha^{-n} \left| h^*(\sigma^n(P)) - \hat{h}_{X,C,L}(\sigma^n(P)) \right| = O(\alpha^{-n}),$$

where the implied constant depends only on X, C, and L. Letting $n \to \infty$, we have $h^* = \hat{h}_{X,C,L}$, as claimed. □

References

[1] P. Autissier, Hauteur des correspondances de Hecke. *Bull. Soc. Math. France* **131** (2003), no. 3, pp. 421–433.

[2] P. Autissier, Pascal, Dynamique des correspondances algbriques et hauteurs. *Int. Math. Res. Not.* (2004), no. 69, pp. 3723–3739.

[3] R. M. Avanzi and U. M. Zannier, Genus one curves defined by separated variable polynomials and a polynomial Pell equation. *Acta Arith.* **99** (2001), no. 3, pp. 227-256.

[4] M. Baker, A Finiteness Theorem for Canonical Heights Attached to Rational Maps over Function Fields. *J. Reine Angew. Math.* **626** (2009), pp. 205–233.

[5] M. Baker and L. DeMarco, Preperiodic Points and Unlikely Intersections, with Laura DeMarco, *Duke Math. J.* **159**, no. 1 (2011), pp. 1–29.

[6] R. L. Benedetto, Heights and preperiodic points of polynomials over function fields. *Inter. Math. Res. Not.* (2005), no. 62, pp. 3855–3866.

[7] Y. F. Bilu and R. F. Tichy, The Diophantine equation $f(x) = g(y)$. *Acta Arith.* **95** (2000), no. 3, pp. 261-288.

[8] G. S. Call and J. H. Silverman, Canonical heights on varieties with morphisms, *Compositio Math.* **89** (1993), pp. 163–205.

[9] A. Chambert-Loir, Algebraic dynamics, function fields and descent, available at `http://www.math.u-psud.fr/~bouscare/workshop_diff/expose1.pdf`

[10] T.-C. Dinh and N. Sibony, Distribution des valeurs de transformations méromorphes et applications, *Comment. Math. Helv.* **81**, no. 1 (2006), pp.221–258.

[11] D. Ghioca and N. M. Mavraki, Variation of the canonical height in a family of rational maps, *New York Journal of Mathematics* **19** (2013), pp. 873–907.

[12] W. Gubler, Heights of subvarieties over M-fields, in *Arithmetic Geometry*, ed. F. Catanese, *Symp. Math.* **37** (1997), pp. 190–227.

[13] M. Hindry and J. H. Silverman, *Diophantine Geometry: An Introduction*, volume 201 of *Graduate Texts in Mathematics*, Springer, New York, 2000.

[14] P. Ingram, Variation of the canonical height for a family of polynomials, *J. Riene Ange. Math.* (to appear).

[15] P. Ingram, Canonical heights for Hénon maps, *Proc. London Math. Soc.* (to appear).

[16] P. Ingram, A lower bound for the canonical height associated to a Drinfeld module, *Int. Math. Res. Notices*, (to appear).
[17] R. Jones, Galois representations from pre-image trees: an arboreal survey. arXiv:1402.6018 (24 Feb 2014)
[18] J. de Jong et al., The Stacks Project, `http://stacks.math.columbia.edu`
[19] S. Lang, *Fundamentals in Diophantine Geometry*, Springer, New York, 1983.
[20] P. Morton and J. H. Silverman, Rational periodic points of rational functions. *Internat. Math. Res. Notices* (1994), no. 2, pp. 97–110.

Fort Collins, USA